

Robust and Opportunistic Autonomous Science for a Potential Titan Aerobot

Daniel M. Gaines, Tara Estlin, Steve Schaffer, Rebecca Castaño and Alberto Elfes

Abstract—

We are developing onboard planning and execution technologies to provide robust and opportunistic mission operations for a potential Titan aerobot. Aerobots have the potential for collecting a vast amount of high priority science data. However, to be effective, an aerobot must address several challenges including communication constraints, extended periods without contact with Earth, uncertain and changing environmental conditions, maneuverability constraints and potentially short-lived science opportunities. We are developing the AerOASIS system to develop and test technology to support autonomous science operations for a potential Titan Aerobot. The planning and execution component of AerOASIS is able to generate mission operations plans that achieve science and engineering objectives while respecting mission and resource constraints as well as adapting the plan to respond to new science opportunities. Our technology leverages prior work on the OASIS system for autonomous rover exploration. In this paper we describe how the OASIS planning component was adapted to address the unique challenges of a Titan Aerobot and we describe a field demonstration of the system with the JPL prototype aerobot.

I. INTRODUCTION

NASA's 2008 Solar System Exploration Roadmap [1] highlights the importance of aerial probes as a strategic new technology for Solar System exploration, and outlines potential missions to Venus and Titan that would use airborne vehicles, such as balloons and airships (blimps). For example, recent data and imagery from Cassini and Huygens show that Titan is a fascinating planetary body with a variety of surface features (See Figure 1). Pictures from the Huygens probe dramatically illustrate the utility of low altitude (< 10 km) aerial imagery at Titan, showing river channels and other striking terrain features not visible from orbit. Airborne vehicles provide a promising means of exploring these areas as they offer traversal capabilities, geographical coverage, and speeds that are orders of magnitude greater than rovers, leading to an enormous science data collection potential.

However, bandwidth constraints, communications latencies and blackouts, and flight maneuvering limitations require these vehicles to have onboard autonomous science capabilities. This would allow science data acquisition to be planned




Fig. 1. Examples of geographic features of interest on Titan.




Fig. 2. Autonomous science can provide the ability for an aerobot to respond to events before they are passed by.

and executed in real time, and would consequently maximize the mission science return. Specifically, autonomous science technology would provide the ability to 1) prioritize data for downlink such that most important data is downlinked early, 2) summarize and compress data, for example, by compiling and downlinking statistics on terrain features observed rather than complete image sets, and 3) detect and respond to science opportunities onboard the vehicle before they are passed over (Figure 2).

We are developing the Aerial Onboard Autonomous Science Investigation System (AerOASIS) to provide onboard science capabilities for aerial probes. Figure 3 illustrates the concept. The illustration shows a Titan aerobot (robotic




Fig. 3. An aerobot science operations plan showing goals that made it into the plan and goals in reserve for potential replanning.

blimp) concept surveying the Huygens descent region at 6 km altitude. The AerOASIS software would analyze the data obtained from onboard multi-spectral imagers and other sensors, and would use scientist-defined signatures to prioritize data, search for high-value science targets and plan science activities. In this example, AerOASIS has selected three science sites: a possible methane lake, a drainage system, and some distant dunes. The aerobot would plan its flight trajectory taking into account local wind conditions, the topography, and its own flight maneuverability and power limitations, and perform close-up surveys of the three selected sites.

In this paper we describe the design of the AerOASIS system with an emphasis on how AerOASIS provides robust and opportunistic planning and execution to support autonomous science on an aerobot. Goal for the aerobot may originate from ground operations on Earth as well as from the vehicle itself by onboard science analysis algorithms. The planning and execution system develops a mission plan that attempts to maximize the value of the goals accomplished while respecting resource and mission constraints.

One of the significant challenges faced by an aerial vehicle is handling the large degree of uncertainty in the environment. The system must be robust to unexpected events that result in deviations from the expected course of events. For example, changes in wind conditions may result in longer transit times and increased energy consumption to reach certain goals. The system must be opportunistic to take response to newly detected science objectives as well as to take exploit fortuitous events, such as favorable wind conditions reducing the expected time to reach a goal. Using the CASPER continuous planning and executing system, we have developed autonomous technology that enables an aerobot to generate mission operation plans and adjust the plans to accommodate unexpected events during execution (Figure 4). This enables the system to appropriately respond to unexpected events and to take advantage of new science




Fig. 4. An aerobot science operations plan showing goals that made it into the plan and goals in reserve for potential replanning.

opportunities.

The paper will also describe a field demonstration of the planning and execution system of AerOASIS that was conducted with the JPL prototype aerobot. It demonstrated AerOASIS' ability to generate and execute valid mission plans and to adapt the plan to unexpected events including new science opportunities.

II. AEROASIS: AERIAL ONBOARD AUTONOMOUS SCIENCE INVESTIGATION SYSTEM

The Aerial Onboard Autonomous Science Investigation System (AerOASIS) (Figure 5) allows an airborne planetary exploration vehicle to:

- 1) summarize and prioritize the most scientifically relevant data from the various incoming sensor streams for relay to Earth;
- 2) identify and select high-value science sites for additional investigation by the aerial vehicle (through low-altitude, high resolution surveys; in-situ probe deployment; and/or surface sample acquisition); and
- 3) dynamically plan, schedule and monitor the various science activities being performed by the aerial vehicle, even during extended communications blackout periods with Earth.

The AerOASIS system is composed of three main subsystems:

- 1) Feature Extraction, which processes sensor imagery and other types of data (such as atmospheric pressure, temperature, wind speeds, etc.) and performs data segmentation and feature extraction.
- 2) Data Analysis and Prioritization, which matches the extracted feature vectors against scientist-defined signatures. The results are used to 1) detect novelty, i.e. statistically significant new types of information or data correlations; 2) perform science data prioritization and summarization for downlink to Earth; 3) identify and select high-value science sites for in-situ studies to be conducted by the aerial vehicle.
- 3) Planning and scheduling, which generates operations plans to achieve observation requests submitted from Earth and from onboard data analysis. These science

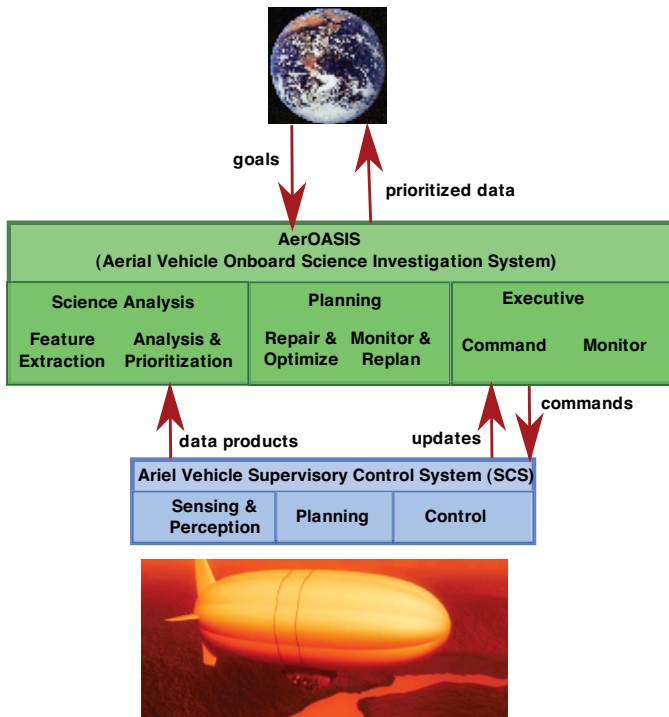


Fig. 5. The AerOASIS Aerial Vehicle Autonomous Science System.

requests can include low-altitude high-resolution surveys, in-situ sonde deployment, and/or surface sample acquisition for onboard analysis.

The AerOASIS system receives images and other sensor data streams from the aerial vehicle, and performs data segmentation and feature extraction. The resulting feature vectors are processed by the data analysis and prioritization subsystem, which performs the identification and selection of high-value science sites, as well as the prioritization and summarization of science data for downlink to Earth. The selected high-priority science sites, as well as additional science requests, are handed as observation requests to the planning and scheduling component which attempts to accomplish these new requests. The planner interfaces with the Aerial Vehicle Supervisory Control System (SCS) to execute the plan and to receive updates on the current state of the vehicle and the world.

The Aerial Vehicle Supervisory Control System (SCS) supervises all sensing/perception, planning, flight navigation and control activities of the aerial exploration vehicle and its deployable sensors and probes [2], [3]. SCS provides the lower levels of the onboard autonomy architecture, including sensor and actuator control, vehicle state estimation, power management, the Flight Control System (FCS), and the underlying flight mode controller, as well as intermediate levels of autonomy such as navigation and flight planning, image-based motion estimation (IBME) for GPS-denied vehicle localization, image mosaicking and geographical mapping, and 3D terrain structure estimation.

In the remainder of this paper, we discuss the planning and scheduling component of AerOASIS in more detail.

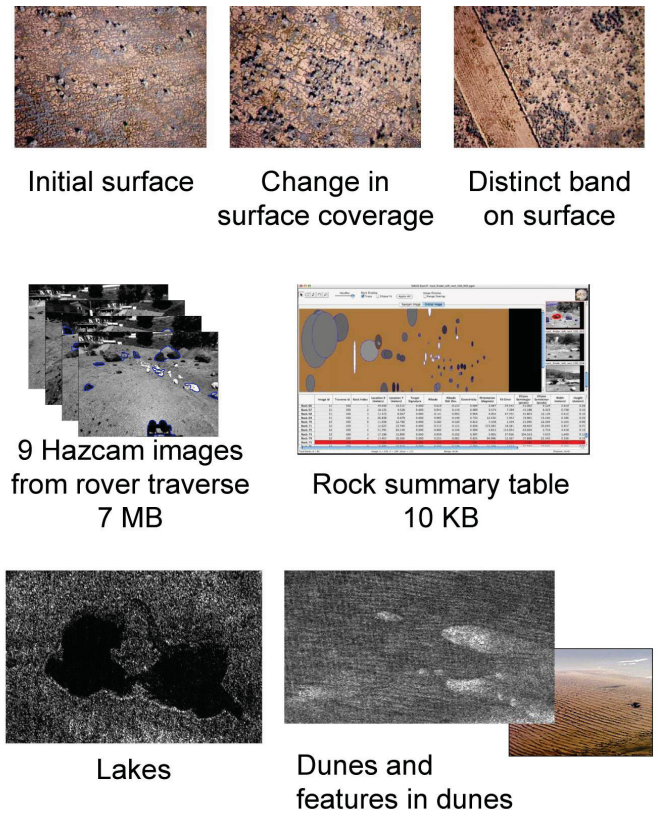


Fig. 6. Science analysis in the AerOASIS system.

III. AUTONOMOUS SCIENCE ANALYSIS THE AEROASIS SYSTEM

AerOASIS performs three major types of science analysis (Figure 6).

Data Prioritization: The mobility and data gathering capabilities of an aerobot coupled with the communications constraints anticipated for a potential Titan mission mean that an aerobot's ability to gather data will far outpace its ability to downlink that data. Ideally, the most interesting data would be downlinked first with less interesting data saved for later. Toward this end, AerOASIS performs prioritization to rank acquired data. There are a variety of techniques that can be employed including: target signature analysis, to look for data that contains objects that match a scientist-provided specification; novelty detection, to look for data that contains unusual information; and characteristic analysis, to prioritize data that serves as a representative sample of a larger population of data [4].

For Titan image analysis, we anticipate that texture analysis will be an important factor in evaluating science quality of images. Texture can be used to differentiate different types of terrain such as plains and dunes. Detecting a change in terrain as the aerobot travels can be indicative of a boundary between terrain types which is of particular scientific interest.

Data Symmetrization: Rather than sending full, raw data products to Earth, AerOASIS can generate highly compact summaries of the contents of data. Figure 6 shows an exam-

ple from a rover collecting images during navigation. Images are processed to identify rocks. For each identified rock, a set of features is extracted (location, albedo, size, shape, ...). A table is generated that summarizes the rocks identified during the traverse. Downlinking the table rather than all the images results in a dramatic reduction in downlink volume. On Earth, a coarse representation of the area can be reproduced from the table. Additionally, thumbnails of the identified rocks can be downlinked if additional imagery is desired to augment the table.

Event Detection: Science analysis is also used to identify potentially interesting scientific events such as the detection of terrain features (e.g. lakes, dunes, ...) or, as mentioned previously, the transition between different terrain types. When an event is detected, the science analysis unit sends a goal request to the planner to collect additional observations. These goals are represented the same as goals from Earth and, like the goals from Earth, include priority information to help the planner decide whether or not to include the goal into its resource-constrained plan.

IV. PLANNING AND EXECUTION IN THE AEROASIS SYSTEM

Our objective is to enable onboard planning software to generate correct and high quality operations plans to achieve mission objectives issued from ground operations as well as respond to science opportunities detected onboard the vehicle. In particular, the system considers prioritized observation requests. This will enable the ground team to uplink a larger set of observations and let the aerobot dynamically select among them based on the scientific and engineering merit of the resulting plan and the aerobot's assessment of available resources. During execution, the aerobot will modify the plan based on the current estimate of its resources.

Our approach is implemented within the CASPER system [5], [6]. CASPER employs a continuous planning technique where the planner continually evaluates the current plan and modifies it when necessary based on new state and resource information. Rather than consider planning a batch process, where planning is performed once for a certain time period and set of goals, the planner has a current goal set, a current aerobot state, and state projections into the future for that plan. At any time an incremental update to the goals or current state may update the current plan. This update may be an unexpected event (such as a new science target) or a current reading for a particular resource level (such as battery charge). The planner is then responsible for maintaining a plan consistent with the most current information.

A plan consists of a set of grounded (i.e., time-tagged) activities that represent different aerobot actions and behaviors. Aerobot state in CASPER is modeled by a set of plan timelines, which contain information on states, such as aerobot position, and resources, such as energy. Timelines are calculated by reasoning about activity effects and represent the past, current and expected state of the aerobot over time. As time progresses, the actual state of the aerobot drifts from the state expected by the timelines, reflecting changes

in the world. If an update results in a problem, such as an activity consuming more memory than expected and thereby over-subscribing RAM, CASPER re-plans, using iterative repair [7], to address conflicts.

CASPER includes an optimization framework for reasoning about soft constraints such as reducing the distance traversed by the aerobot and increasing the value of science data collected. User-defined preferences are used to compute plan quality based on how well the plan satisfies these constraints. Optimization proceeds similar to iterative repair. For each preference, an optimization heuristic generates modifications that could potentially improve the plan score.

Figure 7 provides a high level description of the control algorithm used for the aerobot application of CASPER. The algorithm takes as input a set of goals with associated science priorities and a set of time and resource constraints. CASPER's optimization framework supports a wide-range of user-defined preferences. The main loop of the algorithm interleaves iterative repair and iterative optimization to search for a conflict-free plan of high quality. The loop begins by processing any updates on state and resource timelines or on activity status. It then enters a loop in which it attempts to improve the plan by repairing conflicts or performing optimization steps.

```

Input
  A set of science observations (oversubscribed)
  Time, resource constraints & Preferences

Repeat:
  Process updates from Executive
  Optimize for  $n$  iterations
    If no conflicts
      Select a preference to improve
      Select and apply an improvement method
    Else
      Select conflict to work on
      Select and apply a repair strategy
  Compute plan score
  If current plan is best seen so far, save it
  Reload plan with highest score
  Commit and/or Rescind activities
  If idle, attempt to move up future activities

```

Fig. 7. CASPER control algorithm for aerobot domain.

If there are no conflicts, CASPER attempts to improve the plan by increasing the score of one of the preferences (e.g. by satisfying an observation to increase one of the science quality metrics). If there are conflicts, it will perform an iteration of repair, selecting one of the available repair methods (e.g. move an activity, add an activity, ...). If deletion of an observation is selected, it will select the observation that contributes least to the preference that was most recently selected for improvement.

Note that satisfying an observation will likely introduce conflicts as this is where CASPER will evaluate the resource and temporal requirements of an observation. CASPER will

use subsequent iterations to try to resolve these conflicts. For example, if the aerobot is not currently at the appropriate location to take an observation, CASPER will identify a state conflict which it will attempt to resolve. One option for fixing this conflict is to add an activity that could move the aerobot from one location to another, i.e. a traverse activity. This is also where CASPER selects an ordering of observations in an attempt to minimize traverse distance. We use a simple traveling salesman heuristic to pick start times for activities to reduce traverse distance.

Figure 8 illustrates the “lifetime” of observations in the system. New observations are placed in a requested bin. When an observation is selected to be satisfied, it moves from requested to pending in which it awaits execution. In the meantime, it may be deleted to resolve conflicts in the plan, in which case it moves back to requested. As it nears time for a pending observation to be executed, it is committed and sent to an executive process for execution. If a problem occurs in the plan before the actual execution time of the activity, the planner has the ability to request a rescind of the observation from the executive. If the executive is able to honor the rescind request, it is as if the observation had been deleted from the plan and it returns to the requested bin.




Fig. 8. The lifetime of an observation.

AerOASIS uses the Task Description Language (TDL) for its executive [8]. TDL provides an interface between the planner and the underlying flight architecture. TDL decomposes activities from the planner into specific commands for the aerobot. TDL also performs tighter, closed-loop monitor of task execution that is feasible within the planner’s algorithm. For example, Given the potentially high rate of travel of the aerobot, close monitoring of position is needed for accurate imaging. Thus, within AerOASIS, TDL is responsible for timing of imaging to make sure images are acquired at the correct location. TDL receives sensory data from the aerobot avionics and reports back to CASPER with updates on the state of the world and vehicle resources. TDL also provides updates on activity status including informing CASPER if activities complete early or are running late. This allows CASPER to propagate the affects of these duration changes and adjust the plan if appropriate.



Fig. 9. The JPL Aerobot.

V. FIELD DEMONSTRATION

The JPL Aerobot (Figure 9) is a robotic airship, i.e., a self-propelled lighter-than-air (LTA) vehicle developed by the JPL Aerobot team. The airship specifications are as follows: length of 11 m, diameter of 2.5 m, total volume of 34 m³, maximum speed of 13 m/s (25 kts), maximum ceiling of 3000 m (although, for safety reasons, it is flown below 1000 m), average mission endurance of 60 minutes, static lift payload of 12 kg ASL (at sea level), and dynamic lift payload of up to 16 kg ASL. The avionics and communication systems are installed in the gondola. The forward and aft ropes in Fig. 7 (a,b) are mooring lines for ground handling, not tether lines; i.e., the Aerobot flies free, not tethered.

The aerobot avionics is built around a PC-104+ computer architecture. The navigation sensors consist of an IMU (angular rates, linear accelerations), inclinometers (yaw, roll and pitch angles), magnetic compass, laser altimeter (surface relative altitude), barometric altimeter (absolute altitude against reference point), differential GPS (absolute 3D position) and ultrasonic anemometer (relative wind speed/direction).

The imaging sensors currently consist of a down-looking navigation camera and a wide-angle science camera. The navigation camera is used for image-based motion estimation (IBME), i.e., the ability to estimate the aerial vehicle trajectory in a GPS-denied environment through registration of sequences of images. The results are fed into an extended Kalman filter, together with other sensor estimates, to provide motion and positions estimates for the Aerobot. Other vision-based products include image mosaicking for mapping of large regions, and motion-based estimates of 3D terrain




Fig. 10. JPL Aerobot flying during field trials at the Southern California Logistics Airport (SCLA) in Victorville, CA.

structure. The science camera is used for science image acquisition, science target selection, and visual go-to target and stationkeep over target capabilities using a visual servoing approach to flight control. An onboard control switching system allows toggling between autonomous flight control and human pilot control. For safety reasons, the human pilot always has ability to override the SCS system and reassert control over the Aerobot.

In November, 2008, we successfully demonstrated this technology on the JPL aerobot platform at the Southern California Logistics Airport (SCLA) in Victorville, CA (Figure 10). The flight began with two imaging observation requests at different latitude/longitude coordinates submitted to the system from the “ground team.” The planner generated an operations plan to visit these locations after determining that there was sufficient time and energy to do so.

The winds were relatively strong and, for one of the observations, the aerobot reached the goal much sooner than initially predicted due to tail winds. Since the planner continually monitors the current state, this early arrival was detected and the image request was sent down for execution in time. The second goal required the aerobot to fly against the wind resulting in longer traverse time than the first goal. Again, the planner monitored the situation and waited until the aerobot came close to the target location before commanding the imaging.

While the aerobot was traveling toward the second goal, we simulated an onboard opportunistic science request,

causing a new imaging goal to be sent to the planner. The planner correctly added the request to the plan and achieved this new goal after completing the second goal.

The flight demonstrated the system’s ability to generate mission operations plans given science goals, to command the aerobot control software and monitor plan execution, and to successfully modify the operations plan to respond to dynamic, opportunistic science events.

VI. CONCLUSIONS

The AerOASIS system provides autonomous planning and execution capabilities for aerial vehicles. The system is capable of generating high quality operations plans that integrate observation requests from ground planning teams as well as opportunistic science events detected onboard the vehicle while respecting mission and resource constraints. We have successfully demonstrated the planning component in a field demonstration integrated with the JPL Aerobot platform. In future work we plan to extend the field demonstration with integrated, onboard science analysis.

ACKNOWLEDGMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. This work was funded by the Jet Propulsion Laboratory Research and Technology Development program.

REFERENCES

- [1] NASA, “2008 solar system exploration roadmap,” NASA, Washington, DC, Tech. Rep., 2008.
- [2] A. Elfes, J. L. Hall, J. F. Montgomery, C. F. Bergh, and B. A. Dudik, “Towards a substantially autonomous aerobot for exploration of titan,” in *Proceedings of IEEE International Conference on Robotics and Automation*, Las Vegas, May 2004.
- [3] A. Elfes, J. Hall, E. Kulczycki, A. Morfopoulos, J. Montgomery, D. Clouse, D. Bayard, J. Cameron, R. Machuzak, L. Magnone, C. Bergh, and G. Walsh, “Autonomy capability development for a titan aerobot,” in *Proceedings of the Fourth Annual International Planetary Probe Workshop (IPPW-4)*, Pasadena, California, June 2006.
- [4] R. Castano, T. Estlin, M. Judd, D. Gaines, A. Castano, B. Bornstein, T. Stough, K. Wagstaff, and R. Anderson, “Opportunistic science with a rover traverse science data analysis system,” in *Proceedings of the 8th International Symposium on Artificial Intelligence, Robotics and Automation in Space*, Munchen, Germany, 2005.
- [5] T. Estlin, F. Fisher, D. Gaines, C. Chouinard, S. Schaffer, and I. Nenas, “Continuous planning and execution for a Mars rover,” in *Third International NASA Workshop on Planning and Scheduling for Space*, Houston, TX, October 2002.
- [6] S. Chien, R. Knight, A. Stechert, R. Sherwood, and G. Rabideau, “Using iterative repair to improve the responsiveness of planning and scheduling,” in *Fifth International Conference on Artificial Intelligence Planning and Scheduling*, Breckenridge, CO, April 2000.
- [7] M. Zweben, B. Daun, E. Davis, and M. Deale, “Scheduling and rescheduling with iterative repair,” in *Intelligent Scheduling*, M. Fox and M. Zweben, Eds. Morgan Kaufmann Publishers Inc., 1994, pp. 241–256.
- [8] Simmons and Apfelbaum, “A task description language for robot control,” in *Proceedings of Conference on Intelligent Robotics and Systems*, Vancouver, Canada, October 1998.